

to interact with any of a plurality of disparate 3P agents that enable performance of actions across a plurality of disparate service sectors. For example, the common automated assistant interface may be utilized to engage any one of a plurality of 3P agents that perform intended action(s) in a “restaurant reservation” service sector, engage any one of a plurality of 3P agents that perform intended action(s) in a “purchasing professional services” service sector, engage any one of a plurality of 3P agents that perform intended action(s) in a “purchasing travel services” service sector, and engage any one of a plurality of 3P agents that perform intended action(s) in an “interactive game” service sector.

As used herein, a 3P agent references one or more computing devices and/or associated software managed by a party that is separate from a party that manages an automated assistant. The 3P agent is configured to receive (e.g., over a network via an API) content from the automated assistant. In response to receiving the content, the 3P agent generates 3P responsive content based on the received content, and transmits the 3P responsive content for the provision of output that is based on the 3P responsive content. For example, the 3P agent may transmit the 3P responsive content to the automated assistant for provision of output, by the automated assistant, that is based on the 3P responsive content. A 3P agent may often be configured to perform one or more particularized intended actions such as, for example: booking a restaurant reservation; ordering food; purchasing movie tickets; purchasing services; requesting services (e.g., transportation); managing text, email, or other electronic communications of a user; providing guidance for a task of a user (e.g., mixing a drink, fixing a leaky faucet); engaging in an interactive game with a user; etc.

Various types of input are described herein that may be provided by a user, via user interface input device(s), to an automated assistant and/or to a 3P agent. In some instances the input may be natural language input that is free-form, such as textual input that is based on user interface input generated by the user via one or more user interface input devices (e.g., based on typed input provided via a physical or virtual keyboard or based on spoken input provided via a microphone). As used herein, free-form input is input that is formulated by a user and that is not constrained to a group of options presented for selection by the user (e.g., not constrained to a group of options presented in a drop-down menu).

In some implementations, a method performed by one or more processors is provided and includes: receiving natural language input of a user that is directed to an automated assistant as part of a dynamic dialog between the user and the automated assistant. The natural language input is free-form and based on user interface input provided by the user via a user interface input device. The method further includes selecting an intended action based on determining that the natural language input indicates the intended action. The intended action is selected from a group of potential intended actions that are available via the automated assistant and that span multiple service sectors. The method further includes determining that at least one mandatory parameter stored in association with the intended action is unspecified in the natural language input. The method further includes, in response to determining that the at least one mandatory parameter is unspecified, generating a prompt that is based on the at least one mandatory parameter. The method further includes: providing the prompt as part of the dynamic dialog and as a reply to the natural language input; receiving additional natural language input provided by the

user as part of the dynamic dialog in response to the prompt; determining a value for the mandatory parameter based on the additional natural language input; and selecting a particular third-party agent from a group of third-party agents that can each perform the intended action indicated by the natural language input. The method further includes transmitting a third-party invocation request that includes the value for the mandatory parameter. The transmitting is to the particular third-party agent via one or more network interfaces and occurs without any other third-party invocation request transmission to any other third-party agents of the group of third-party agents. In some implementations, the method further includes: receiving, via one or more of the network interfaces, responsive content from the third-party agent in response to the transmitting the third-party invocation request; and providing, for presentation to the user as part of the dynamic dialog, output that is based on the responsive content.

These and other implementations of technology disclosed herein may optionally include one or more of the following features.

In some implementations, the method further includes: receiving further input provided by the user in response to the output that is based on the responsive content of the third-party agent; and transmitting, to the third-party agent via one or more of the network interfaces, content that is based on the further input. In some of those implementations, the further input is voice input and the method further includes converting the voice input to text and including at least some of the text in the content. In some versions of those implementations, the method further includes: receiving, via one or more of the network interfaces, further responsive content from the third-party agent in response to transmitting the content; and providing, for presentation to the user as part of the dynamic dialog, further output that is based on the further responsive content.

In some implementations, the method further includes: receiving further input provided by the user in response to the output that is based on the responsive content of the third-party agent; determining whether the further input indicates a desire to cease interaction with the third-party agent; and in response to determining that the further input does not indicate a desire to cease interaction with the third-party agent: transmitting, to the third-party agent via one or more of the network interfaces, content that is based on the further input.

In some implementations, the method further includes: receiving further input provided by the user in response to either the output or further output that is based on content provided by the third-party agent; determining that the further input indicates a desire to interact with another third-party agent; and in response to determining that the further input indicates a desire to interact with another third-party agent: transmitting, to an alternative third-party agent of the group of third party agents, an additional third-party invocation request that includes the value for the mandatory parameter. In some of those implementations, the method further includes: determining an additional value for an additional parameter of the intended action based on user input during engagement of the third-party agent in the dynamic dialog; and including the additional value for the additional parameter in the additional third-party invocation request.

In some implementations, the method further includes: determining an additional value for a non-solicited optional parameter of the intended action based on the natural language input or the additional natural language input; and